arXiv:1609.03686v1 [stat.ME] 13 Sep 2016

New multivariate tests for assessing covariate balance in matched observational studies

BY HAO CHEN

Department of Statistics, University of California, Davis
One Shields Ave., Davis, CA 95616 USA
hxchen@ucdavis.edu

AND DYLAN S. SMALL

Department of Statistics, The Wharton School, University of Pennsylvania,
3730 Walnut Street, Philadelphia, PA 19104 USA
dsmall@wharton.upenn.edu

SUMMARY

We propose new tests for assessing whether covariates in a treatment group and matched control group are balanced in observational studies. The tests exhibit high power under a wide range of multivariate alternatives, some of which existing tests have little power for. The asymptotic permutation null distributions of the proposed tests are studied and the p-values calculated through the asymptotic results work well in finite samples, facilitating the application of the test to large data sets. The tests are illustrated in a study of the effect of smoking on blood lead levels. The proposed tests are implemented in an `R` package `BalanceCheck`.

Some key words: Holistic test; Minimum spanning tree; Nearest neighbor; Non-parametric test; Permutation null distribution.

1. INTRODUCTION

In testing the effect of treatments, random assignment is desirable because it tends to make the treated group and the control group similar in both observed and unobserved covariates. However, in many applications, for practical and ethical reasons, we may not be able to assign subjects randomly. For example, if we want to study the effect of smoking, we could not randomly pick some people and force them to smoke and force some other people not to smoke. Rather, we must conduct an observational study.

In a matched observational study, we hope to measure all the confounders and then construct a treated (e.g., smoker) group and a matched control (e.g., not smoker) group that are similar with respect to these confounders. If there is only one confounder, it may be feasible to match treated units exactly to control units, but there are often many confounders in which case matching exactly is infeasible. For example, if there are 20 binary confounders, there are 2^{20} possible values of the confounders. So even with thousands of subjects, it will typically be impossible to match a treated subject exactly to a control subject. Although typically we cannot find exact matches when there are many confounders, we could still hope that the distribution of the covariates is similar in the treated and matched control group – this is what a randomized experiment hopes to achieve, not to create identical treated and control groups.

A key question for a matched observational study is, has the matching created a treated group and a matched control group that have similar distributions? If the answer is yes, then any substantial differences in outcomes between the two groups can be attributed to the treatment while if the answer is no, then a difference in outcomes between the two groups may just be due to pre-treatment differences between the groups. A commonly used approach for comparing the treated and matched control groups is to compare the distribution of the measured confounders of the two groups (e.g., gender, age, height, weight, profession, etc.) one by one, for example by conducting two sample t-tests. A problem with this univariate approach is that the joint distributions of the covariates might be quite different between the two groups even though their marginal distributions are similar.

Some more holistic tests have been suggested. Rosenbaum & Rubin (1985) proposed to use the propensity score, the probability of a subject being assigned to a particular treatment given a set of covariates, to check covariate balance. By summarizing all observed covariates into one value, it greatly simplified the process and has been applied in many studies since its launch (Berk & Newton, 1985; Hoffer et al., 1985; Myers et al., 1987; Fiebach et al., 1990; Czajka et al., 1992; Stone et al., 1995; Lieberman et al., 1996). However, to estimate the propensity score, one needs to model the distribution of the treatment indicator variable given the observed covariates. So choosing a good model is important for this approach. More recently, some less model-based approaches have been proposed and used in several studies. Heller et al. (2010) proposed to use the crossmatch test of Rosenbaum (2005) to test covariate balance through forming a non-bipartite graph among the treatment and control subjects. Distributions that are very different between the treated and control groups will exhibit few cross-matches (matches between treated and control subjects). Hansen & Bowers (2008) advocated a randomization inference cousin of Hotelling's T^2 statistic for balance checking.

When the number of covariates is large, the behavior of "closeness" could be counterintuitive due to the curse of dimensionality. The existing methods more or less use the implicit assumption that "closer" observations are more likely to be from the same distribution. This assumption may not hold when the dimension of the observation (number of covariates) is high (Chen & Friedman, 2016). In this paper, we propose new tests to address this problem in covariate balance assessment. The proposed tests still use similarity information among the observations, while the test statistics are designed so that more alternative scenarios could be covered under typical sample sizes, as well as taking into account the effect of matching. The proposed tests are implemented in an `R` package `BalanceCheck`.

The rest of the paper is organized as follows. Section 2 reviews existing two-sample tests based on similarity information. Section 3 discusses the proposed tests and their asymptotic properties. The power of the proposed tests is evaluated and compared to existing tests in Section 4. The proposed tests are applied to an observational study on the effect of smoking on blood lead levels in Section 5. The paper concludes with discussion in Section 6.

2. A BRIEF REVIEW ON TWO-SAMPLE TEST BASED ON SIMILARITY INFORMATION

The setup of the general two-sample testing is as follows: There are two independent samples $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ and $\{\mathbf{y}_1, \ldots, \mathbf{y}_m\}$, where the units in each sample are independent and identically distributed according to distributions $F_{\mathbf{x}}$ and $F_{\mathbf{y}}$, respectively. The setup of covariate balance assessment is slightly different as the two samples are no long independent due to the matching step, while some ideas could still be adopted. We here review the existing two-sample tests based on similarity information as the proposed tests build up on them.

Friedman & Rafsky (1979) proposed the first practical two-sample test for multivariate data based on similarity information. They used the pairwise distances among the pooled observations from the two samples to construct a minimal spanning tree (MST), which is a spanning tree that connects all observations with the sum of distances of edges in the tree minimized. The test statistic is the number of edges that connect nodes (observations) from different samples and the null hypothesis of equal distribution is rejected if this number of significantly *less* than its null expectation. The rationale of the test is that, if the two samples are from the same distribution, the two samples are well mixed and we expect a relative large number of edges that connect subjects from different samples. So if the number of between-sample edges is small, the two samples are from different distributions. We call the test the *edge-count test* for easy reference.

The edge-count test was later extended to other similarity graphs, such as nearest neighbor (NN) graphs in Schilling (1986) and Henze (1988), and a non-bipartite matching graph in Rosenbaum (2005) (crossmatch). Maa et al. (1996) showed that the edge-count test based on MST constructed on Euclidean distance is consistent against all alternatives for multivariate data.

Nevertheless, in a recent study, Chen & Friedman (2016) found that the edge-count test has low or even no power under some common alternatives in practical sample sizes. For example, when the two distributions are $F_{\mathbf{x}} = \mathcal{N}(\mathbf{0}, I_d)$ and $F_{\mathbf{y}} = \mathcal{N}(\boldsymbol{\mu}, \sigma^2 I_d)$, $\|\boldsymbol{\mu}\|_2 = 1, \sigma = 1.1, d = 100, n = m = 1,000$, the number of between-sample edges in the MST constructed on the pooled 2,000 observations is close to its null expectation in typical runs. In contrast, if $\sigma = 1$ and all other parameters are the same, the number of between-sample edges is significantly smaller than its null expectation in typical runs. This is a weird phenomenon – the MST test has more power when the distributions are less different. The underlying reason for the above weird phenomenon is the curse of dimensionality: The volume of the d-dimensional ball grows exponentially in dimension, so the observations from the sample with a slightly higher variance tends to find the observations from the sample with a slightly smaller variance to be closer, resulting in the large number of between-sample edges.

Chen & Friedman (2016) solved this problem by proposing a modified test statistic based on the following fact: In the first scenario above where the edge-count tests work, which we call 'Scenario (a)', the number of within-sample edges for both samples are significantly *larger* than their null expectations. In the second scenario above where the edge-count test does not work, which we call 'Scenario (b)', the number of within-sample edges for the sample with slightly smaller variance is significantly *larger* than its null expectation, while the number of within-sample edges for the sample with slightly larger variance is significantly *smaller* than its null expectation. The test proposed in Chen & Friedman (2016) aggregates deviations in both directions of the within-sample edges for the two samples from their null expectations, and thus works for more alternatives than the tests that only rely on between-sample edges. We call this test the *generalized edge-count test* for easy reference.

The edge-count test has been applied to assess covariate balance in several studies (Heller et al., 2010; McHugh et al., 2013). However, the generalized edge-count test can lead to wrong conclusions under the setting of observational studies. Consider the scenario that the matching is done perfectly. If we construct a MST on the observations, most edges would be connections between the two groups, i.e., both numbers of within-sample edges would be small. The test in Chen & Friedman (2016) also rejects in this scenario as these numbers also significantly deviate from their null expectations. This issue does not exist for testing two random sample because perfect matching for two randomly selected samples happens with extremely low probability. However, this is common in matching, especially when the matching is done well.

In the following, we propose tests that have substantial power under both scenarios (a) and (b), while tending not to reject in the scenario that the matching balances the covariates.

3. NEW TESTS

Let n be the number of treated subjects (subjects exposed to a treatment, such as smoking) and $\mathbf{y}_i,\ i = 1, \ldots, n$, be the vector of observed covariates for subject i. There are more than n controls and among them n subjects are matched to the treated subjects according to some rules. Let $\mathbf{y}_{n+1}, \ldots, \mathbf{y}_{2n}$ be the observed covariate vectors of the matched controls. Let $N = 2n$ be the total sample size. We test whether the two sets of covariates are similar through testing whether the group identities, treatment and control, is exchangeable. In other words, we work under the permutation null distribution that place $1/\binom{N}{n}$ probability on each of the $\binom{N}{n}$ choices of n out of N observations as the treatment group. When there is no further specification, we denote by $\mathbf{P}$, $\mathbf{E}$, $\mathbf{Var}$, probability, expectation, and variance, respectively, under the permutation null distribution.

We propose two new tests with one based on the nearest neighbor (NN) information of the subjects (Section 3·1) and the other based on the MST constructed on the pooled subjects (Section 3·2). We call the tests "CrossNN" and "CrossMST", respectively.

3·1. *CrossNN*

We pool the $N = 2n$ observations together and for each observation, finds its NN among the pooled N observations. Let $a_{ij} = 1$ if $\mathbf{y}_j$ is the NN of $\mathbf{Y}_i$ and 0 if otherwise. Let D_{12} be the number of observations in treatment group whose NNs are from the controls, and D_{21} be the number of observations in controls whose NNs are from the treatment group. For any event x, let $\mathbf{I}_x$ be the indicator function that equals 1 if event x occurs and 0 if otherwise. Then, D_{12} and D_{21} can be expressed as

$$D_{12} = \sum_{i=1}^{n} \sum_{j=n+1}^{N} a_{ij} = \sum_{i=1}^{N} \sum_{j=1}^{N} a_{ij}\, \mathbf{I}_{i\in[1,n]}\, \mathbf{I}_{j\in[n+1,N]}, \tag{1}$$

$$D_{12} = \sum_{i=1}^{n} \sum_{j=n+1}^{N} a_{ji} = \sum_{i=1}^{N} \sum_{j=1}^{N} a_{ji}\, \mathbf{I}_{i\in[1,n]}\, \mathbf{I}_{j\in[n+1,N]}. \tag{2}$$

We propose to use $D_m = \min(D_{12}, D_{21})$ as the test statistic. We reject the null of balanced covariates if D_m is significantly *smaller* than its expectation under the permutation null distribution. The statistic defined in this way has power for scenarios (a) and (b) because at least one of D_{12} and D_{21} is small under the alternatives of either scenario. On the other hand, under the perfect matching scenario, both D_{12} and D_{21} are large and the null of balanced covariates is not rejected.

The next question is how small D_m needs to be to reject the null of covariate balance. When n is small, we can permute the group label directly to obtain the permutation distribution of D_m and compute the permutation p-value. However, when n is large, direct permutation would be too time consuming. We thus seek to find an analytic approach to do so.

In the following, we first derive exact analytic expressions for the expectations and variances of D_{12} and D_{21}, as well as their covariance. We then study the joint distribution of D_{12} and D_{21} asymptotically, which equips us to determine the asymptotic p-value for the test. The asymptotic p-value is then compared to the p-value calculated through 10,000 permutations directly. It turned out that the asymptotic p-value is reasonably accurate for sample sizes in hundreds (see Figure 1).

LEMMA 1. *The expectations and variances of D_{12} and D_{21}, as well as their covariance are as follows:*

$$\mathbf{E}(D_{12}) = \mathbf{E}(D_{21}) = \frac{N^2}{4(N-1)} \triangleq \mu_D, \tag{3}$$

$$\begin{aligned}\mathbf{Var}(D_{12}) &= \mathbf{Var}(D_{21}) \\ &= \frac{1}{16}\left(\frac{N^2}{N-1} - \frac{N^2}{(N-1)^2} + 2C_1\frac{N(N-2)}{(N-1)(N-3)} + 2C_2\frac{N(N-4)}{(N-1)(N-3)}\right) \triangleq \sigma_D^2,\end{aligned} \tag{4}$$

$$\mathbf{Cov}(D_{12}, D_{21}) = \frac{1}{16}\left(\frac{N^2}{N-1} - \frac{N^2}{(N-1)^2} + (2C_1 - 2C_2)\frac{N(N-2)}{(N-1)(N-3)}\right), \tag{5}$$

where C_1 is the number of pairs of observations that are mutual NNs and C_2 is the number of pairs of observations sharing their NNs.

The lemma is proved through combinatorial analysis and the details are in Appendix A·1.

Remark 1. According to Propositions 1 and 2 in (Henze, 1988), when $N \to \infty$, we have $2C_1/N$ and $2C_2/N$ converge to constants that only depend on the dimension and the norm used to calculate the distance. We denote the two limits by q_1 and q_2, respectively.

Let

$$Z_{D,1} = \frac{D_{12} - \mu_D}{\sigma_D}, \quad Z_{D,2} = \frac{D_{21} - \mu_D}{\sigma_D}.$$

Then the test statistic $Z_D = \min(Z_{D,1}, Z_{D,2})$ is equivalent to D_m. The following theorem states the asymptotic joint distribution of $Z_{D,1}$ and $Z_{D,2}$.

THEOREM 1. *Under permutation null distribution, as $N \to \infty$, $(Z_{D,1}, Z_{D,2})'$ converges to a bivariate Gaussian distribution with mean $\mathbf{0}$ and covariance matrix $\begin{bmatrix} 1 & \rho_D \\ \rho_D & 1 \end{bmatrix}$, where $\rho_D = (1 + q_1 - q_2)/(1 + q_1 + q_2)$.*

The proof of this Theorem extends the methods in Chen & Zhang (2013) and Chen & Friedman (2016). The complete proof is in Appendix A·2.

Base on this Theorem, we can easily calculate the asymptotic p-value of Z_D. That is,

$$P(Z_D \le z) = 1 - P(Z_D > z) = 1 - P(Z_{D,1} > z, Z_{D,2} > z),$$

and $P(Z_{D,1} > z, Z_{D,2} > z)$ can be calculated from function `pmvnorm()` in `R` package `mvtnorm`. In practice, we use the finite sample version of the covariance between $Z_{D,1}$ and $Z_{D,2}$, that is, we use

$$\rho_{D,N} = \frac{N(N-3)/(N-1) + 2C_1 + 2C_2(N-4)/(N-2)}{N(N-3)/(N-1) + 2C_1 - 2C_2}$$

rather than ρ_D when calculating $P(Z_{D,1} > z, Z_{D,2} > z)$. It is easy to see that $\lim_{N\to\infty} \rho_{D,N} = \rho_D$.

We compared the asymptotic p-values to p-values calculated from 10,000 permutations. Figure 1 plots the boxplots of the differences between the two p-values (asymptotic p-value minus permutation p-value) over 100 simulation runs under difference choices of n and under different dimensions. We see that when the dimension is low ($d = 10$), the asymptotic p-value is already quite accurate for very small sample sizes. The accuracy for very small sample sizes is not as

good when the dimension becomes high ($d = 100$), but it is reasonably accurate once the sample size is at least 100.

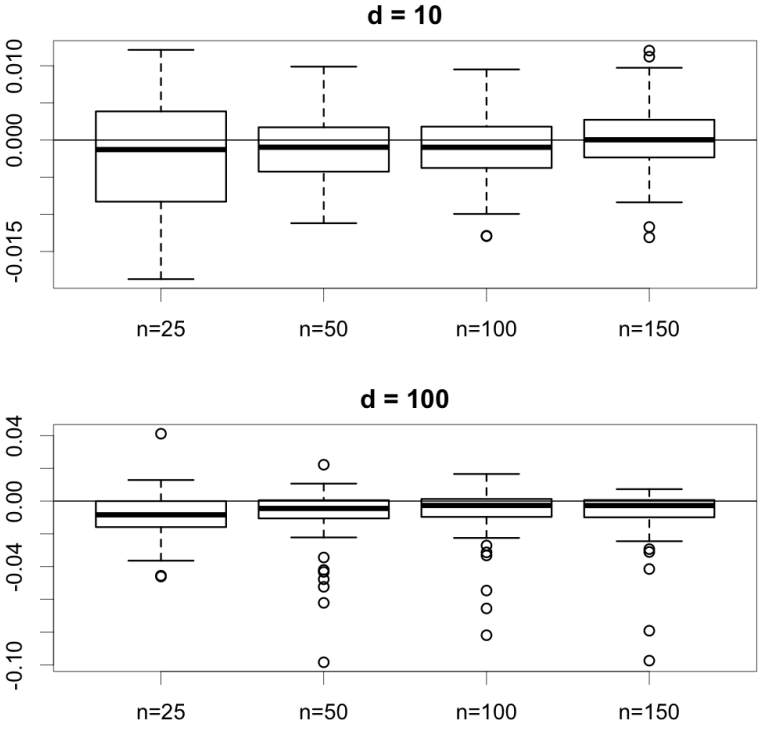

Fig. 1. Boxplots of the differences between the asymptotic p-value and the p-value calculated from 10,000 permutations (asymptotic p-value minus permutation p-value).

3·2. *CrossMST*

We pool the n treated subjects and n matched controls together and construct a MST on the pooled observations. Let R_1 be the number of edges in the MST that connect subjects from the treatment group and R_2 be the number of edges that connect subjects from the control group. We propose to use $R_M = \max(R_1, R_2)$ as the test statistic. We reject the null of balanced covariates if R_M is significantly *larger* than its expectation under the permutation null distribution.

The statistic defined in this way has power for scenarios (a) and (b) because at least one of R_1 and R_2 is large under the alternatives of each scenario. On the other hand, under the perfect matching scenario, both R_1 and R_2 are small and the null of balanced covariates is not rejected. Similar to D_m, we also studied the permutation null distribution of R_M and the p-value approximation based on asymptotic results is also reasonable accurate for sample sizes in hundreds. The details for these are in Appendix B.

4. POWER ASSESSMENT

We now evaluate the power of the proposed tests, CrossNN and CrossMST, in comparison to three existing tests: Hotelling's T^2 test, the test proposed in Hansen & Bowers (2008), and the crossmatch test proposed in Rosenbaum (2005). We follow a similar simulation setting to that in Franklin et al. (2014).

The detailed setup is as follows: There are 6 observed covariates that are independent standard normal distributed $(X_1, X_2, X_3, X_4, X_5, X_6)$. The exposure/treatment, T, depends on these 6 covariates as well as the second order of four of them $(X_7 = X_1^2, X_8 = X_2^2, X_9 = X_3^2, X_{10} = X_4^2)$. In particular, T was simulated as a binary variable via the logistic model $\text{logit}\{\text{P}(T_i = 1)\} = \alpha_0 + \boldsymbol{\alpha}\mathbf{X}_i$, where $\mathbf{X}_i$ is the covariate vector (including all 10 covariates) for subject i and $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_{10})$. The parameters in $\boldsymbol{\alpha}$ define the log-odds ratios between covariates and exposure in the pre-matched dataset, and higher absolute values generally indicate more imbalance. We generated 1,000 subjects and determined whether each of them is exposed or not. Then, we matched control subjects $(T = 0)$ to exposed subjects $(T = 1)$ through their propensity scores on $X_1, \ldots, X_6$, the main effects of the covariates.

To make the comparison simple, we let $\alpha_1 = \cdots = \alpha_6 = a$ and $\alpha_7 = \cdots = \alpha_{10} = b$. We considered three scenarios: (i) there are only main effects $(a = 0.4, b = 0)$; (ii) there are both main and second order effects $(a = b = 0.4)$; and (iii) there are only second order effects $(a = 0, b = 0.4)$. Since the propensity score is calculated based on the main effects, we would expect the matching could balance the covariates when there are only some main effects. However, if there are second order effects, the matching would not be able to address them.

For each scenario, we simulated 100 data sets and calculated the standardized difference between the treatment group and matched control group for each of the ten covariates for each data set. Standardized difference is defined as $= (\bar{x}_1 - \bar{x}_2)/\sqrt{(s_1^2 + s_2^2)/2}$, where $\bar{x}_m$ and s_m^2 are sample mean and variance for treated subjects $(m = 1)$ and matched controls $(m = 2)$. Figure 2 plots the boxplots of standardized differences of the 100 data sets for each of the ten covariates. We see that the main effects $(X_1, \ldots, X_6)$ are well balanced in all three scenarios. So the matching over propensity scores is doing a good job for covariates included in the model. However, second order effects $(X_7, \ldots, X_{10})$ are very unbalanced except for scenario (i). This is expected as they are not included in the model for calculating the propensity score and only scenario (i) does not have second order effects from the data generating step.

Hence, we want a test that will not reject covariate balance for scenario (i), while reject covariate balance for scenarios (ii) and (iii). Table 1 shows the proportion of trials, out of 100, that the test reject covariate balance at the significance level 0.05 for the three scenarios. We see that all tests are doing a good job for scenario (i) that the proportions of rejections are all controlled by the significance level. For scenarios (ii) and (iii), we want a high rejection proportion. We see that the Hotelling's T^2 and Hansen's test do not have power in these two scenarios. The other three tests have power in these two scenarios. The CrossMST has higher power than the CrossNN test and both have higher power than the crossmatch test. (The values of a and b were chosen so that these three tests have moderate power to be comparable.)

Table 1. *The proportion of trials (out of 100) that the test rejects covariate balance at the significance level 0.05.*

Scenario	CrossNN	CrossMST	Hotelling's T^2	Hansen's test	crossmatch
(i)	0.05	0.02	0	0	0.03
(ii)	0.47	0.74	0	0	0.17
(iii)	0.38	0.79	0	0	0.21

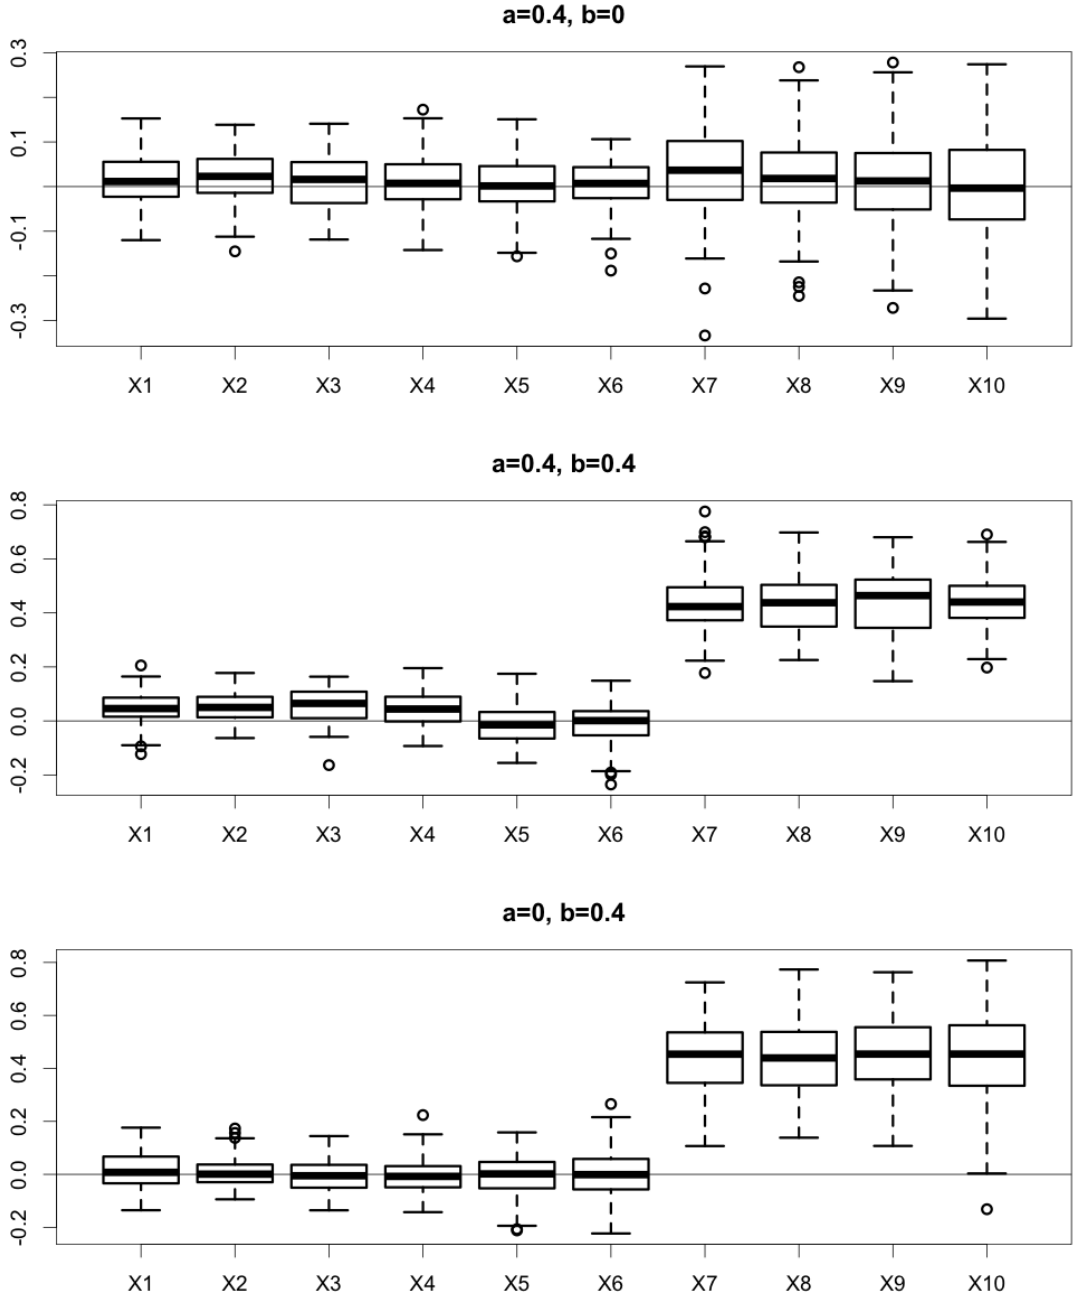

Fig. 2. Boxplots of standardized differences between the treatment group and matched control group over 100 trials.

5. APPLICATION TO A STUDY OF SMOKING'S EFFECT ON BLOOD LEAD

In Hsu & Small (2013), the authors studied the effect of smoking on blood lead levels among 679 daily smokers and 2661 non-smokers using data from the 2007 to 2008 U.S. National Health and Nutrition Examination Survey (NHANES). A daily smoker reported smoking every day for

the previous 30 days and smoking an average of at least 10 cigarettes per day on these days. A non-smoker reported smoking fewer than 100 cigarettes in his or her life and smoking no cigarettes in the previous 30 days. All subjects were at least 20 years old and had no tobacco use besides cigarette smoking in the previous 5 days. The question is whether smoking causes an increase in blood lead levels.

We first matched 679 daily smokers with 679 non-smokers, among a candidate pool of 2661 non-smokers, with the hope that the covariates in the control group are balanced with the exposed/treated group. We matched using the `pairmatch()` function in the `optmatch` package in `R` (Hansen, 2007) on a distance matrix that combined a caliper on an estimated propensity score with a rank based Mahalanobis distance (Rosenbaum, 2010). Pairs were matched for age, gender, education, income, and race. Table 2 shows means and standardized differences in means (i.e., the difference in means divided by a measure of the average within treatment group standard deviation) before and after matching.

Table 2. *Summary statistics and covariates balance before and after matching for pairs of a daily smoker and a non-smoker.*

Covariates	Before			After	
	Smoker ($n = 679$)	Non-smoker ($n = 2661$)	SD*	Non-smoker ($n = 679$)	SD*
Age	46.71	49.78	-0.188	46.79	-0.005
Income-to-poverty level	2.00	2.62	-0.412	2.15	-0.098
Missing,%	6.04	9.17	-0.118	6.77	-0.028
Male,%	56.85	37.81	0.388	56.11	0.015
Education,%					
Less than 9th grade	11.19	13.83	-0.080	9.13	0.062
9-11th grade	26.51	13.90	0.318	23.42	0.078
High school graduate	32.25	23.26	0.202	33.73	-0.033
Some college	24.74	25.63	-0.020	26.80	-0.048
College	5.30	23.30	-0.532	6.92	-0.048
Unknown	0.00	0.08	-0.039	0.00	0.000
Race,%					
White	65.24	40.81	0.505	63.48	0.037
Black	19.15	19.13	0.000	21.06	-0.049
Mexican American	6.48	21.35	-0.440	7.07	-0.017
Other Hispanic	4.71	13.38	-0.306	3.83	0.031
Other races	4.42	5.34	-0.043	4.57	-0.007

*SD: Standardized difference $= (\bar{x}_1 - \bar{x}_2)/\sqrt{(s_1^2 + s_2^2)/2}$, where $\bar{x}_m$ and s_m^2 are sample mean and variance for smokers $(m = 1)$ and non-smokers $(m = 2)$.

We can see from the table that the covariates are quite unbalanced before matching. For example, some covariates have absolute standardized difference larger than 0.5. After doing the matching, the covariates seem to be balanced well. All absolute standardized differences are less than 0.1. We applied the standard two-sample t-test on each of the covariates after matching. The lowest p-value is larger than 0.05. Hence, the covariates are well balanced separately.

Now the question is whether the covariates are well balanced jointly. We applied the Hotelling T^2 test to the matched data and get a p-value of 0.69, showing no significant difference. We also

applied the crossmatch test to the matched data and get a p-value of 0.014, showing some evidence of a difference between the two groups. We then applied our proposed methods, CrossNN and CrossMST, to the matched data and get p-values of 0.00034 and 0.00077, respectively, which provides very strong evidence rejecting covariate balance.

To see which of these discrepant test results seems to better assess the evidence for covariate imbalance, we explored the data set more. We regressed the smoking status on the covariates (1) with no interaction term, and (2) with pairwise interactions between covariates. Categorical variable `education` has six categories, and none of the treated subjects and matched controls is in the category "Unknown", so we use four dummy variables to represent `education` in the model. Categorical variable `race` has five categories and is represented by four dummy variables in the model. This leads to 12 variables in model (1) and 53 additional interaction terms in model (2).

When no interaction term was included in the regression, no coefficient in the regression model is significantly different from 0 except for the intercept. This is in line with the fact that all covariates are quite well balanced separately. However, when all pairwise interaction terms are included, quite a few coefficients become significantly different from 0. They are listed in Table 3. We see that, after including the interaction terms, the variables `age` and `education` exhibit main effects and there are also extensive interaction effects between `age` and `education`.

Table 3. *Coefficient estimates and standard errors for terms whose coefficients are significantly different 0 (p-value less than 0.01).*

Covariates	Estimate	Standard Error	t-value	p-value
Age	0.0206	0.0063	3.268	0.00111
Education (Less than 9th grade)	1.0613	0.3791	2.799	0.00520
Education (9-11th grade)	0.9355	0.3464	2.700	0.00701
Age $\times$ Education (Less than 9th grade)	-0.0184	0.0047	-3.924	9.17×10^{-5}
Age $\times$ Education (9-11th grade)	-0.0189	0.0040	-4.696	2.94×10^{-6}
Age $\times$ Education (High school graduate)	-0.0160	0.0039	-4.076	4.86×10^{-5}
Male $\times$ Education (9-11th grade)	0.3627	0.1406	2.580	0.00999

Therefore, the covariates of the two groups are not jointly balanced. Our new tests report very small p-values, giving an appropriate alarm whereas the existing tests provided less of an alarm.

6. CONCLUSION AND DISCUSSION

We propose two tests, CrossNN and CrossMST, to assess covariate balance in matched observational studies. Both tests utilize similarity information on the covariates and are powerful for a wide range of alternatives, some of which existing tests have little power for.

Both CrossNN and CrossMST can be extended to denser graphs. For example, CrossNN can be extended to k-NN where each subject finds its first k nearest neighbors. Similarly, CrossMST can be extended to k-MST, which is the union of 1, ..., kth MST, where ith MST is a spanning tree connecting all observations that minimizes the sum of distances across edges subject to the constraint that this spanning tree does not contain any edge in the 1st, ..., (i-1)th MST(s). Choosing k slightly larger than 1 could further boost the power of the tests as the slightly denser graphs contains more similarity information.

In addition, since both tests only need a similarity measure on the sample space, it is not limited to multivariate data. For example, the observation on each covariate does not need to be a value. It could be a curve over time such as the blood pressure over a month. As long as a reasonable distance can be defined over the observations, the proposed tests could be applied.

ACKNOWLEDGEMENT

Hao Chen is supported in part by NSF award DMS-1513653.

A. PROOFS TO LEMMAS AND THEOREMS

A·1. *Proof to Lemma* 1

We work under the permutation null distribution, i.e., the indices $\{1,\dots,N\}$ are permuted. In the following, without further specification, the summation is from 1 to N. Let $\pi(i)$ be the index of $\mathbf{y}_i$ under permutation. We have

$$\begin{aligned}
\mathbf{E}(D_{12}) &= \mathbf{E}\left(\sum_{i,j} a_{ij}\,\mathbf{I}_{\pi(i)\in[1,n]}\,\mathbf{I}_{\pi(j)\in[n+1,N]}\right)\\
&= \sum_{i,j} a_{ij}\,\mathbf{P}(\pi(i)\in[1,n],\pi(j)\in[n+1,N]) = \sum_{i,j} a_{ij}\frac{n^2}{N(N-1)} = \frac{N^2}{4(N-1)},\\
\mathbf{E}(D_{12}^2) &= \mathbf{E}\left(\sum_{i,j,l,r} a_{ij}\,a_{lr}\,\mathbf{I}_{\pi(i)\in[1,n]}\,\mathbf{I}_{\pi(j)\in[n+1,N]}\,\mathbf{I}_{\pi(l)\in[1,n]}\,\mathbf{I}_{\pi(r)\in[n+1,N]}\right)\\
&= \sum_{i,j,l,r} a_{ij}\,a_{lr}\,\mathbf{P}\left(\pi(i)\in[1,n],\pi(j)\in[n+1,N],\pi(l)\in[1,n],\pi(r)\in[n+1,N]\right)\\
&= \sum_{i,j} a_{ij}\,\mathbf{P}\left(\pi(i)\in[1,n],\pi(j)\in[n+1,N]\right)\\
&\quad + \sum_{i,j,l} a_{ij}\,a_{lj}\,\mathbf{P}\left(\pi(i)\in[1,n],\pi(j)\in[n+1,N],\pi(l)\in[1,n]\right)\\
&\quad + \sum_{i,j,l,r \text{ all different}} a_{ij}\,a_{lr}\,\mathbf{P}\left(\pi(i)\in[1,n],\pi(j)\in[n+1,N],\pi(l)\in[1,n],\pi(r)\in[n+1,N]\right)\\
&= N\frac{n^2}{N(N-1)} + \frac{n^2(n-1)}{N(N-1)(N-2)}\sum_{i,j,l,l\neq j} a_{ij}\,a_{lj}\\
&\quad + \frac{n^2(n-1)^2}{N(N-1)(N-2)(N-3)}\sum_{i,j,l,r \text{ all different}} a_{ij}\,a_{lr}\\
&= \frac{N^2}{4(N-1)} + \frac{N}{8(N-1)}2C_2 + \frac{N(N-2)}{16(N-1)(N-3)}(N^2-3N+2C_1-2C_2).
\end{aligned}$$

Then, $\mathbf{Var}(D_{12})$ follows from $\mathbf{Var}(D_{12}) = \mathbf{E}(D_{12}^2) - (\mathbf{E}(D_{12}))^2$. Following the same argument, we get the expectation and variance for D_{12}.

For the covariance between D_{12} and D_{21}, we have

$$\mathbf{E}(D_{12}D_{21}) = \mathbf{E}\left(\sum_{i,j,l,r} a_{ij}\,a_{rl}\,\mathbf{I}_{\pi(i)\in[1,n]}\,\mathbf{I}_{\pi(j)\in[n+1,N]}\,\mathbf{I}_{\pi(l)\in[1,n]}\,\mathbf{I}_{\pi(r)\in[n+1,N]}\right)$$

$$
\begin{aligned}
&= \sum_{i,j,l,r} a_{ij}\, a_{rl}\, \mathbf{P}\left(\pi(i) \in [1,n], \pi(j) \in [n+1,N], \pi(l) \in [1,n], \pi(r) \in [n+1,N]\right) \\
&= \sum_{i,j} a_{ij}\, a_{ji}\, \mathbf{P}\left(\pi(i) \in [1,n], \pi(j) \in [n+1,N]\right) \\
&\quad + \sum_{i,j,r,j\neq r} a_{ij}\, a_{ri}\, \mathbf{P}\left(\pi(i) \in [1,n], \pi(j) \in [n+1,N], \pi(r) \in [n+1,N]\right) \\
&\quad + \sum_{i,j,l,i\neq l} a_{ij}\, a_{jl}\, \mathbf{P}\left(\pi(i) \in [1,n], \pi(j) \in [n+1,N], \pi(r) \in [n+1,N]\right) \\
&\quad + \sum_{i,j,l,r \text{ all different}} a_{ij}\, a_{rl}\, \mathbf{P}\left(\pi(i) \in [1,n], \pi(j) \in [n+1,N], \pi(l) \in [1,n], \pi(r) \in [n+1,N]\right) \\
&= 2C_1 \frac{n^2}{N(N-1)} + \frac{n^2(n-1)}{N(N-1)(N-2)} 2(N-2C_1) \\
&\quad + \frac{n^2(n-1)^2}{N(N-1)(N-2)(N-3)}(N^2 - 3N + 2C_1 - 2C_2).
\end{aligned}
$$

Then the covariance of D_{12} and D_{21} follows from $\mathbf{Cov}(D_{12}, D_{21}) = \mathbf{E}(D_{12}D_{21}) - \mathbf{E}(D_{12})\mathbf{E}(D_{21})$.

A·2. *Proof to Theorem* 1

The proof of Theorem 1 relies on Stein's method. Consider sums of the form $W = \sum_{i\in\mathcal{J}} \xi_i$, where $\mathcal{J}$ is an index set and ξ are random variables with $\mathbf{E}\xi_i = 0$, and $\mathbf{E}(W^2) = 1$. The following assumption restricts the dependence between $\{\xi_i : i \in \mathcal{J}\}$.

*Assumption A.*1. (Chen & Shao, 2005, p. 17) For each $i \in \mathcal{J}$ there exists $K_i \subset L_i \subset \mathcal{J}$ such that ξ_i is independent of $\xi_{K_i^c}$ and ξ_{K_i} is independent of $\xi_{L_i^c}$.

We will use the following existing theorem in proving Theorem 1.

THEOREM A.1. *(Chen & Shao, 2005, Theorem 3.4) Under Assumption A.*1*, we have*

$$\sup_{h\in Lip(1)} |E(h(W)) - E(h(Z))| \leq \delta$$

where $Lip(1) = \{h : \mathbb{R} \to \mathbb{R}\}$*,* Z

$$\delta = 2\sum_{i\in\mathcal{J}} (E|\xi_i\eta_i\theta_i| + |E(\xi_i\eta_i)|E|\theta_i|) + \sum_{i\in\mathcal{J}} E|\xi_i\eta_i^2|$$

with $\eta_i = \sum_{j\in K_i} \xi_j$ *and* $\theta_i = \sum_{j\in L_i} \xi_j$*, where* K_i *and* L_i *are defined in Assumption A.*1.

To prove the theorem, we take one step back by studying (D_{12}, D_{21}) under the bootstrap null distribution, which is defined as follows: For each observation, we assign it to be from the treatment group with probability 0.5. Let n_B be the number of observations that are assigned to be treated subjects. Then, conditioning on $n_B = n$, the bootstrap null distribution becomes the permutation null distribution. We use $\mathbf{P}_B, \mathbf{E}_B, \mathbf{Var}_B$ to denote the probability, expectation, and variance under the bootstrap null distribution, respectively.

Under the bootstrap null distribution, it is not hard to show that

$$
\begin{aligned}
\mathbf{E}_B(D_{12}) &= \frac{N}{4} \\
\mathbf{E}_B(D_{12}^2) &= \frac{1}{16}(N^2 + N + 2C_1 + 2C_2)
\end{aligned}
$$

Let

$$\mu_D^B = \frac{N}{4}, \quad (\sigma_D^B)^2 = \frac{1}{16}(N + 2C_1 + 2C_2),$$

then $E_B(D_{12}) = E_B(D_{21}) = \mu_D^B$, $Var_B(D_{12}) = Var_B(D_{21}) = (\sigma_D^B)^2$.

Let

$$W_1^B = \frac{D_{12} - \mu_D^B}{\sigma_D^B}, \qquad W_2^B = \frac{D_{21} - \mu_D^B}{\sigma_D^B}, \qquad W_3^B = \frac{(n_B - n)}{\sqrt{N/4}}.$$

We next show that, when $N \to \infty$,

(a) $(W_1^B, W_2^B, W_3^B)'$ becomes multivariate Gaussian distributed under the bootstrap null.
(b)

$$\frac{\sigma_D^B}{\sigma_D} \to 1, \quad \frac{\mu_D^B - \mu}{\sigma_D^B} \to 0.$$

From (a) and given that $\mathbf{Var}_B(W_3^B) = 1$, we have that $(W_1^B, W_2^B)'|W_3^B$ follows bivariate Gaussian distribution under the bootstrap null distribution as $N \to \infty$. Since the permutation null distribution is equivalent to the bootstrap null distribution given $W_3^B = 0$, $(W_1^B, W_2^B)'$ follows bivariate Gaussian distribution under the permutation null distribution as $N \to \infty$. Since

$$Z_{D,1} = \frac{D_{12} - \mu_D}{\sigma_D} = \frac{\sigma_D^B}{\sigma_D}\left(W_1^B + \frac{\mu_D^B - \mu_D}{\sigma_D^B}\right), \quad Z_{D,2} = \frac{D_{21} - \mu_D}{\sigma_D} = \frac{\sigma_D^B}{\sigma_D}\left(W_2^B + \frac{\mu_D^B - \mu_D}{\sigma_D^B}\right),$$

given (b), we have $(Z_{D,1}, Z_{D,2})'$ follow bivariate Gaussian distribution under the permutation null distribution as $N \to \infty$.

To prove (a), by Cramér-Wold device, we only need to show that $W = a_1 W_1^B + a_2 W_2^B + a_3 W_3^B$ is Gaussian distributed for any combination of a_1, a_2, a_3 such that $Var_B(W) > 0$.

Let $G = \{e = (i, j) : \mathbf{y}_j \text{ is the nearest neighbor of } \mathbf{y}_i\}$, and

$$\begin{aligned} A_e =& \{e\} \cup \{e' \in G : e' \text{ and } e \text{ share a node}\}, \\ B_e =& A_e \cup \{e'' \in G : \exists e' \in A_e, \text{ such that } e'' \text{ and } e' \text{ share a node}\}. \end{aligned}$$

So A_e is the subgraph in G that connects to edge e, and B_e is the subgraph in G that connects to any edge in A_e. Let $|G|$ denote the number of edges in the graph G.

For observation i, let $T_i = 1$ if it is assigned to be in the treatment group under the bootstrap distribution, and $T_i = 2$ if it is assigned to be in the control group. For $e = (e^+, e^-)$, where e^- is the nearest neighbor of e^+, let $T_e = 3$ if e^+ is assigned to be in the treatment group and e^- is assigned to be in the control group, and $T_e = 4$ if the other way around.

Let

$$\begin{aligned} \xi_e &= a_1 \frac{I_{T_e=3} - 1/4}{\sigma_D^B} + a_2 \frac{I_{T_e=4} - 1/4}{\sigma_D^B}, \\ \xi_i &= a_3 \frac{I_{T_i=1} - 1/2}{\sqrt{N/4}}. \end{aligned}$$

Then

$$W = \sum_{e \in G} \xi_e + \sum_i \xi_i.$$

Let $a = \max(|a_1|, |a_2|, |a_3|)$, we have $\xi_i \le a/\sqrt{N}$, $|\xi_e| \le 3a/\sqrt{N}$.
Let $\mathcal{J} = \{1, \ldots, N\} \cup \{e \in G\}$ the be index set of all edges and nodes.
For $i \in \{1, \ldots, N\}$, let

$$\begin{aligned} K_i &= \{e \in G_i\} \cup \{i\}, \\ L_i &= \{e \in G_{i,2}\} \cup \{\text{nodes in } G_i\}, \end{aligned}$$

where G_i is the subgraph in G that contains edges associated with $\mathbf{y}_i$, and $G_{i,2}$ is the subgraph in G that contains edges associated with at least one node in G_i. Then K_i and L_i satisfy Assumption A.1.

For $e = (e^+, e^-)$, let

$$K_e = A_e \cup \{e^+, e^-\},$$
$$L_e = B_e \cup \{\text{nodes in } A_e\}.$$

Then K_e and L_e satisfy Assumption A.1.

For $j \in \mathcal{J}$, let $\eta_j = \sum_{k\in K_j} \xi_k$, $\theta_j = \sum_{k\in L_j} \xi_k$. By Theorem A.1, we have $\sup_{h\in Lip(1)} |E_B h(W) - \mathbf{E}h(Z)| \leq \delta$ for $Z \sim \mathcal{N}(0,1)$, where

$$\begin{aligned}
\delta &= \frac{1}{Var_B(W)} \left(2\sum_{j\in\mathcal{J}} (E_B|\xi_j\eta_j\theta_j| + |E_B(\xi_j\eta_j)|E_B|\theta_j|) + \sum_{j\in\mathcal{J}} E_B|\xi_j\eta_j^2| \right) \\
&\leq \frac{1}{Var_B(W)} \left(5\sum_{i=1}^{N} \frac{a^3}{N^{3/2}} (|G_i|+1)(|G_{i,2}| + |G_i| + 1) + 5\sum_{e\in G} \frac{27a^3}{N^{3/2}} (|A_e|+2)(|B_e| + |A_e| + 1) \right) \\
&\leq \frac{1}{Var_B(W)} \left(\frac{10a^3}{N^{3/2}} \sum_{i=1}^{N} (|G_i|+1)(|G_{i,2}|+1) + \frac{1215a^3}{N^{3/2}} \sum_{e\in G} |A_e||B_e| \right)
\end{aligned}$$

Notice that for $e = (i,j)$, we have $G_i, G_j \subseteq A_e$, $G_{i,2}, G_{j,2} \subseteq B_e$, so $(|G_i|+1)(|G_{i,2}|+1) \leq (|A_e|+1)(|B_e|+1)$. For each node i, we can randomly pick an edge that has i as one of its end points, then each edge in the graph can be picked at most twice since an edge only has two end points. Therefore,

$$\sum_{i=1}^{N} (|G_i|+1)(|G_{i,2}|+1) \leq 2\sum_{e\in G} (|A_e|+1)(|B_e|+1) \leq 8\sum_{e\in G} |A_e||B_e|.$$

Hence,

$$\delta \leq \frac{1295a^3}{\sqrt{Var_B(W)}} \frac{1}{N^{3/2}} \sum_{e\in G} |A_e||B_e|.$$

From Corollary S1 of Bickel & Breiman (1983), we have $\sup_{1\leq j\leq N} \sum_{i=1}^{N} a_{ij}^+ \leq \mathbb{C}$ where $\mathbb{C}$ is a positive constant. Then, $|A_e| \leq \mathbb{C}+1$ and $|B_e| \leq (\mathbb{C}+1)^2$, so $\sum_{e\in G} |A_e||B_e| = O(|G|) = O(N)$. Hence, $\delta \to 0$ when $N \to \infty$.

Now, we prove result (b).

$$\lim_{N\to\infty} \frac{\sigma_D^B}{\sigma_D} = \lim_{N\to\infty} \frac{\sqrt{N + 2C_1 + 2C_2}}{\sqrt{\frac{N^2}{N-1} - \frac{N^2}{(N-1)^2} + 2C_1\left(\frac{N}{N-1} - \frac{N}{(N-1)(N-3)}\right) + 2C_2\frac{N(N-2)}{(N-1)(N-3)}}} = 1.$$

$$\lim_{N\to\infty} \frac{\mu_D^B - \mu_D}{\sigma_D^B} = -\lim_{N\to\infty} \frac{N/(N-1)}{\sqrt{N + 2C_1 + 2C_2}} = 0.$$

B. PERMUTATION NULL DISTRIBUTION OF THE TEST STATISTIC OF CROSSMST

There are $N-1$ edges in the MST. Following from Chen & Friedman (2016), we have

$$\mathbf{E}(R_1) = \mathbf{E}(R_2) = \frac{N-2}{4} \triangleq \mu_R, \tag{B1}$$

$$\mathbf{Var}(R_1) = \mathbf{Var}(R_2) = \frac{1}{16(N-3)} \left(-(N-2)(N-6) + 2C_3\frac{N(N-4)}{N-1} \right) \triangleq \sigma_R^2, \tag{B2}$$

$$\mathbf{Cov}(R_1, R_2) = \frac{N-2}{16(N-3)} \left(3(N-2) - 2C_3\frac{N}{N-1} \right), \tag{B3}$$

where C_3 is the number of edge pairs that share a common node.

Let

$$Z_{R,1} = \frac{R_1 - \mu_R}{\sigma_R}, \quad Z_{R,2} = \frac{R_2 - \mu_R}{\sigma_R}.$$

Then the test statistic $Z_R = \max(Z_{R,1}, Z_{R,2})$ is equivalent to R_M.

Let G_{MST} be the set of edges in the MST.

$$A_{e,MST} = \{e\} \cup \{e' \in G_{MST} : e' \text{ and } e \text{ share a node}\},$$
$$B_{e,MST} = A_e \cup \{e'' \in G_{MST} : \exists\, e' \in A_{e,MST}, \text{ such that } e'' \text{ and } e' \text{ share a node}\}.$$

Then A_e is the subgraph in G that connects to edge e, and B_e is the subgraph in G that connects to any edge in A_e.

Then, following from Chen & Friedman (2016), we have that, if $\sum_{e\in G} |A_{e,MST}||B_{e,MST}| = o(N^{1.5})$, $C_3 - N = O(N)$, then under permutation null distribution, as $N \to \infty$, $(Z_{R,1}, Z_{R,2})'$ converges to a bivariate Gaussian distribution with mean $\mathbf{0}$ and covariance matrix $\begin{bmatrix} 1 & \rho_R \\ \rho_R & 1 \end{bmatrix}$, where $\rho_R = (3 - 2q_3)/(2q_3 - 1)$ with $q_3 \triangleq \lim_{N\to} C_3/N$.

We then can easily calculate the asymptotic p-value of Z_R. That is,

$$P(Z_R \geq z) = 1 - P(Z_R < z) = 1 - P(Z_{R,1} < z, Z_{R,2} < z),$$

and $P(Z_{R,1} < z, Z_{R,2} < z)$ is calculated using function `pmvnorm()` in `R` package `mvtnorm`. In practice, we use the finite sample version of the covariance between $Z_{R,1}$ and $Z_{R,2}$, that is, we use

$$\rho_{R,N} = \frac{3(N-2) - 2C_3N/(N-1)}{-(N-6) + 2C_3N(N-4)/(N-1)/(N-2)}$$

rather than ρ when calculating $P(Z_{R,1} < z, Z_{R,2} < z)$. It is easy to see that $\lim_{N\to\infty} \rho_{R,N} = \rho$.

We check how the asymptotic p-value approximation works for finite sample sizes. We compare the asymptotic p-value to p-value obtained through 10,000 permutation directly. The boxplots of the differences under different sample sizes are shown in Figure 3. Similarly to CrossNN, the asymptotic p-value approximation works very well for small sample sizes when the dimension is low ($d = 10$). The approximation becomes worse when the dimension becomes high ($d = 100$), while it is reasonably accurate for sample sizes in hundreds.

REFERENCES

BERK, R. A. & NEWTON, P. J. (1985). Does arrest really deter wife battery? an effort to replicate the findings of the minneapolis spouse abuse experiment. *American Sociological Review* , 253–262.

BICKEL, P. J. & BREIMAN, L. (1983). Sums of functions of nearest neighbor distances, moment bounds, limit theorems and a goodness of fit test. *The Annals of Probability* , 185–214.

CHEN, H. & FRIEDMAN, J. H. (2016). A new graph-based two-sample test for multivariate and object data. *Journal of the American statistical association* .

CHEN, H. & ZHANG, N. R. (2013). Graph-based tests for two-sample comparisons of categorical data. *Statistica Sinica* **23**, 1479–1503.

CHEN, L. & SHAO, Q. (2005). Stein's method for normal approximation. *An introduction to Stein's method* **4**, 1–59.

CZAJKA, J. L., HIRABAYASHI, S. M., LITTLE, R. J. & RUBIN, D. B. (1992). Projecting from advance data using propensity modeling: An application to income and tax statistics. *Journal of Business & Economic Statistics* **10**, 117–131.

FIEBACH, N. H., COOK, E. F., LEE, T. H., BRAND, D. A., ROUAN, G. W., WEISBERG, M. & GOLDMAN, L. (1990). Outcomes in patients with myocardial infarction who are initially admitted to stepdown units: data from the multicenter chest pain study. *The American journal of medicine* **89**, 15–20.

FRANKLIN, J. M., RASSEN, J. A., ACKERMANN, D., BARTELS, D. B. & SCHNEEWEISS, S. (2014). Metrics for covariate balance in cohort studies of causal effects. *Statistics in medicine* **33**, 1685–1699.

FRIEDMAN, J. H. & RAFSKY, L. C. (1979). Multivariate generalizations of the Wald-Wolfowitz and Smirnov two-sample tests. *The Annals of Statistics* **7**, 697–717.

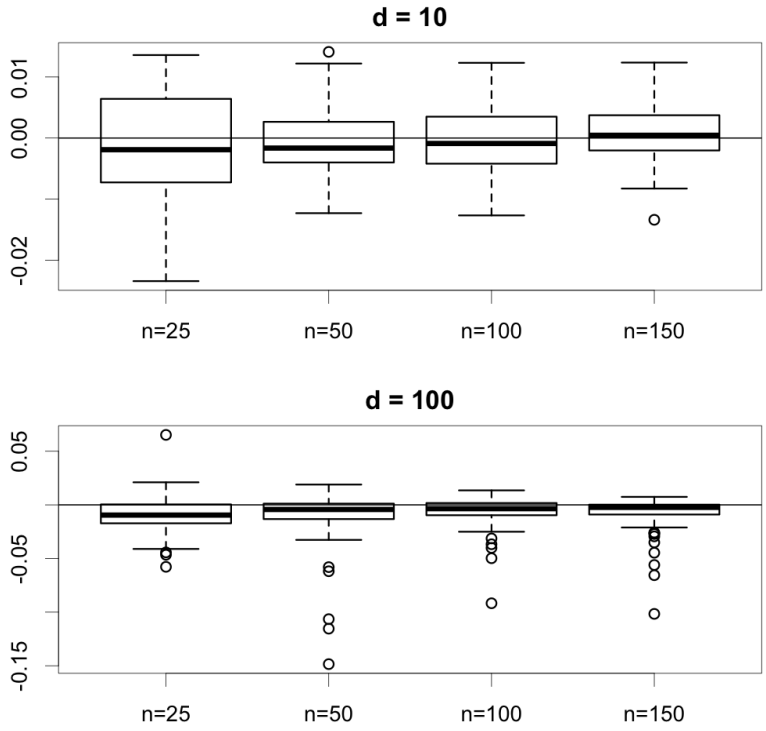

Fig. 3. Boxplots of the differences between the asymptotic p-value and the p-value calculated from 10,000 permutations (asymptotic p-value minus permutation p-value).

HANSEN, B. (2007). Flexible, optimal matching for observational studies. *R News* **7**, 18–24.

HANSEN, B. B. & BOWERS, J. (2008). Covariate balance in simple, stratified and clustered comparative studies. *Statistical Science* , 219–236.

HELLER, R., ROSENBAUM, P. R. & SMALL, D. S. (2010). Using the cross-match test to appraise covariate balance in matched pairs. *The American Statistician* **64**, 299–309.

HENZE, N. (1988). A multivariate two-sample test based on the number of nearest neighbor type coincidences. *The Annals of Statistics* , 772–783.

HOFFER, T., GREELEY, A. M. & COLEMAN, J. S. (1985). Achievement growth in public and catholic schools. *Sociology of Education* , 74–97.

HSU, J. Y. & SMALL, D. S. (2013). Calibrating sensitivity analyses to observed covariates in observational studies. *Biometrics* **69**, 803–811.

LIEBERMAN, E., LANG, J. M., COHEN, A., D'AGOSTINO, R., DATTA, S. & FRIGOLETTO, F. D. (1996). Association of epidural analgesia with cesarean delivery in nulliparas. *Obstetrics & Gynecology* **88**, 993–1000.

MAA, J.-F., PEARL, D. K. & BARTOSZYŃSKI, R. (1996). Reducing multidimensional two-sample data to one-dimensional interpoint comparisons. *The Annals of Statistics* **24**, 1069–1074.

MCHUGH, M. D., BEREZ, J. & SMALL, D. S. (2013). Hospitals with higher nurse staffing had lower odds of readmissions penalties than hospitals with lower staffing. *Health Affairs* **32**, 1740–1747.

MYERS, W. O., GERSH, B. J., FISHER, L. D., MOCK, M. B., HOLMES, D. R., SCHAFF, H. V., GILLISPIE, S., RYAN, T. J. & KAISER, G. C. (1987). Time to first new myocardial infarction in patients with mild angina and three-vessel disease comparing medicine and early surgery: a cass registry study of survival. *The Annals of thoracic surgery* **43**, 599–612.

ROSENBAUM, P. R. (2005). An exact distribution-free test comparing two multivariate distributions based on adjacency. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**, 515–530.

ROSENBAUM, P. R. (2010). Design of observational studies. *Springer series in statistics Show all parts in this series* .

ROSENBAUM, P. R. & RUBIN, D. B. (1985). Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician* **39**, 33–38.

SCHILLING, M. F. (1986). Multivariate two-sample tests based on nearest neighbors. *Journal of the American Statistical Association* **81**, 799–806.

STONE, R. A., OBROSKY, D. S., SINGER, D. E., KAPOOR, W. N., FINE, M. J., TEAM, P. P. O. R. et al. (1995). Propensity score adjustment for pretreatment differences between hospitalized and ambulatory patients with community-acquired pneumonia. *Medical care* , AS56–AS66.